

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
25 January 2001 (25.01.2001)

PCT

(10) International Publication Number  
**WO 01/06798 A1**

(51) International Patent Classification<sup>7</sup>: H04Q 3/545

(21) International Application Number: PCT/US00/18940

(22) International Filing Date: 12 July 2000 (12.07.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
09/356,284 16 July 1999 (16.07.1999) US

(71) Applicant: EXCEL SWITCHING CORPORATION  
[US/US]; 225 Independence Drive, Hyannis, MA 02601  
(US).

(72) Inventor: KRAVITZ, Barry; 29 Thomas Street, Middle-  
boro, MA 02346 (US).

(74) Agents: ROONEY, Rita, M. et al.; Cesari and McKenna,  
L.L.P., 30 Rowes Wharf, Boston, MA 02110 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

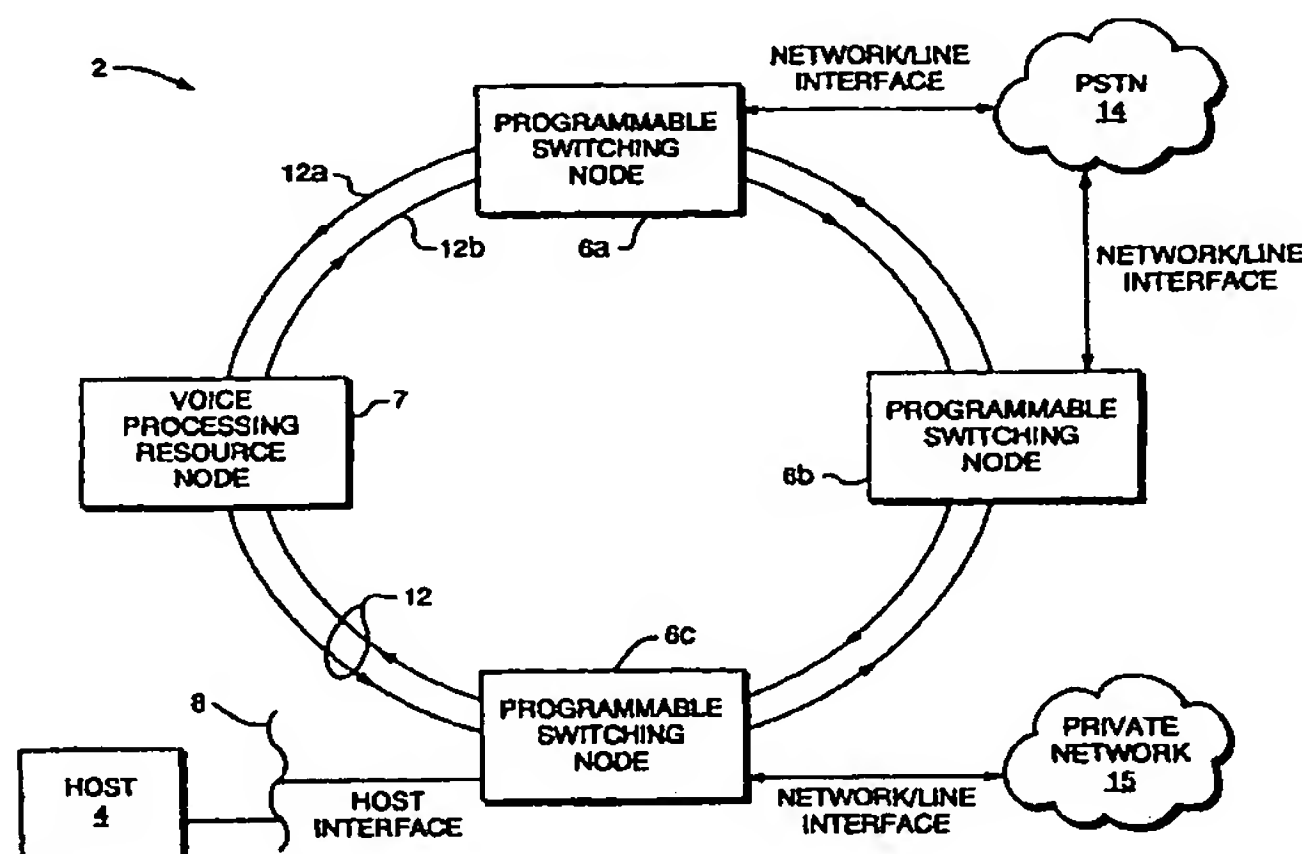
(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— With international search report.

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: METHOD AND APPARATUS FOR IN SERVICE SOFTWARE UPGRADE FOR EXPANDABLE TELECOMMUNICATIONS SYSTEM



(57) Abstract: A method and apparatus is provided for upgrading software, or a portion thereof, on a continuously running system, while the system is operating. A base set of software is built which contains a reserved memory area in its memory architecture. Upgrade source files are built into the reserved memory area. The upgrade images are extracted out and converted into loadable upgrade images. A header indicates where the upgrade is to be downloaded in the system and also identifies the location in the reserved memory area at which the upgrade is to be stored in the relevant component of the system. The upgrade is downloaded into the respective reserved memory area in the memory storage devices in the respective components in the system. The downloads can be performed while the system is operating because the reserved memory area does not contain live code. The upgrades are activated or deactivated using a jump instruction sequence.

WO 01/06798 A1

- 1 -

# **METHOD AND APPARATUS FOR IN SERVICE SOFTWARE UPGRADE FOR EXPANDABLE TELECOMMUNICATIONS SYSTEM**

## **FIELD OF THE INVENTION**

5       The present invention relates generally to upgrading software on a continuously operating computer system, and, more specifically, to effecting such upgrades in the field of telecommunications systems.

## **BACKGROUND OF THE INVENTION**

10       Computer programs which are initially installed and configured on one or more storage devices in the system at start up typically control continuously operating computer systems. The programs are sets of software instructions which perform together to control a variety of functions in many different components of the system.

15       It is often desirable to make modifications to such system software in order to provide additional features to the system, to solve problems or "bugs" which have been found during operation of the system, or to accommodate new developments in technology. Conventionally, when a software upgrade is to be made, a new version of the software code is typically installed and configured on the system. For example, the new code is often installed as a traditional download but for this, normal operation of the running system must be interrupted. Shutting down system operations, in whole or  
20       in part, leads to financial and service losses due to the downtime involved.

25       A mechanism of upgrading stand alone programs, such as word-processing programs or drafting programs, was described in U.S. Patent No. 5,764,992 (Kullick). In accordance with the method described in Kullick, the program checks for an updated or more recent version of itself upon start up. The new version may be resident in the memory on the hard disk or on a network file server. When the current version starts

- 2 -

up and finds that an updated version exists, it overwrites itself with the updated version. The entire program thus replaces itself with a new version during start up, or shut down. Alternatively, the new program can be written to install itself while the computer is not in use. The Kullick patent describes a method of replacing an entire program with a new version of itself while the system is essentially down or not in use.

In a continuously running system, however, such as a telecommunications system, there is no time at which the system is routinely down or not being used. More specifically, a telecommunications system operates continuously as consumers expect to have access to telephone service twenty-four hours a day. Thus, interruptions to system operation can be extremely costly for the service provider and, at the very least, inconvenient to users.

A telecommunications switching system can be comprised of several switching nodes and voice processing resource nodes which are inter-connected by an inter-nodal network, often connected as a ring. Such a system is controlled by software tasks resident on a host computer or a network file server, as well as in microprocessors running on individual cards, like a switching matrix card or a line card and the like. These cards are located within a switching node, voice processing resource node or other application service node. It is often necessary or desirable to upgrade software running on these microprocessors in order to provide additional features, to solve problems which have arisen during the operation of the system or to bring a new card into service. This does not always involve the entire system, but rather, for example, may relate to the software resident on a particular node, or on a portion of software which is running on one type of line card in the system. In addition, many times, a "patch" of a portion of code is to be made by replacing or modifying a portion of the program. In such a case, it is unnecessary to replace all of the software running on a particular type of card, instead it may be that merely a few lines of code are to be replaced.

At present, software upgrades to telecommunications systems and other continuously running systems have been managed as system software releases which are installed as traditional downloads for which system operation must be interrupted, even if only a portion of the system is affected by the upgrade.

- 3 -

A method of replacing an entire program, or a task, running on a telecommunications system has been described. However, this involves a complex mechanism for replacement of substantially the entire task and may also require shut down of certain aspects of the system, or even the entire system.

5        There remains a need, therefore, for a mechanism by which a software upgrade may be effected in a continuously running system, without interruption of operation of the system.

There remains a further need for the ability to amend only a portion of a running software program in such a system or in a subset of the system.

10

### SUMMARY OF THE INVENTION

The present invention is a method and apparatus for performing a live upgrade of the software running on a continuously operating system, such as for example, a telecommunications system, without interrupting the operation of the system.

15        Briefly, in accordance with the invention, a software upgrade mechanism includes a unique method of building a base set of software that includes the capability of accepting an upgrade which can then be activated while the system is on-line. The method of the present invention includes establishing in the base software architecture, a reserved memory area. The reserved memory area is an allotment of a portion of memory that is held available for prospective upgrades; and it represents a correspond-  
20        ing reserved memory area in each memory storage device in each component in the system for which the upgrade is targeted.

In accordance with the invention, when an upgrade is to be built, the first step in the process is to write upgrade source files and place them within the file set of an existing base build. It is assumed that a base build has previously been written. Next, the  
25        upgrade source files are written. The upgrade source files contain various products, such as, for example: code, uninitialized data, initialized data, constants and the like. Each product of a particular type associated with that upgrade is given a predetermined name in accordance with a suitable naming convention.

- 4 -

The base files and the upgrade source files, as well as build instructions are inputs to an appropriate make/build software utility. The build instructions notify an associated compiler and linker where to locate in memory the various portions of the build depending upon the names of those products. The products of the upgrade source files having particular names are located at designated addresses in the reserved memory area. In addition, the compiler and linker resolves the addresses to all referenced global data or other functions in the base software build to ensure that code in the upgrade that calls routines in the base build can identify where those routines are located.

As will be well understood by those skilled in the art, after compiling and linking, a map file is produced. The map file includes information such as the names of the products of the base build and of the upgrade build. It also includes information about where those products are located in memory which ultimately leads to the determination as to where those products will be located in the individual memory storage devices in the system after the download.

In accordance with the invention, a software utility reads the map file and identifies where in the memory image of the overall build the upgrade is located. The utility then extracts the upgrade products and generates a loadable upgrade image which consists of the upgrade (without the base build) and a header.

The header is designed based upon the information in the map file and is associated with the loadable upgrade image. The header may include information about the type of system component, such as a matrix card or a line card in a telecommunications system, for which that particular upgrade is targeted, the address in the reserved memory area of the storage device in that card at which the upgrade is to be stored, as well as other information such as checksum information, date, and the relevant software revision number.

The upgrade is then downloaded to the targeted cards in the system. Such downloads can be performed while the system software is running, and without interruption of the operation of the system, as the reserved memory area does not yet contain live code. The upgrade code that is downloaded into the reserved memory area is later



- 5 -

activated in accordance with the invention. The downloads can thus be performed without overwriting existing code, and without the need to replace the entire program.

After downloading is complete, the upgrade may be activated. All upgrades do not have to be activated at the same time, or at all. It may be that although all line cards  
5 of a particular type receive an upgrade, it is desired to activate the upgrades in only certain of the nodes in the system. In accordance with the invention, the user can determine which upgrades are to be activated.

Activation is performed by overlaying an instruction which transfers control to a different address using a hard jump sequence, for example. In accordance with one as-  
10 pect of the invention, a load-address-and-jump instruction sequence is overlaid at the beginning of an old function, which directs the program to the beginning of a new function, which has been loaded, into the reserved memory area for the upgrade. In an alternative embodiment of the invention, where an address table is being utilized to dynamically call up routines, then the new addresses of the upgraded code are loaded into  
15 available spaces in the jump table which direct the CPU to those new addresses. In either case, the context of the old routine is maintained and the CPU returns to where it would have returned had the old routine been executed. A simple mechanism is used to determine the integrity of activation of upgrades so that a piece of code that is executing at that moment is not patched while executing, as discussed in detail herein.

20 A messaging sequence between the host and a processing card in the system includes a means of querying the status of downloaded upgrades to determine whether that upgrade is activated or deactivated, as well as to compare the upgrades a card actually contains against a host-controlled list of upgrades that such a card type should contain. The query message may also include information about the amount of mem-  
25 ory space available in the reserved memory segments for additional upgrades, if desired.

- 6 -

### BRIEF DESCRIPTION OF THE DRAWINGS

The above and further advantages of the invention may be better understood by reference to the following description in conjunction with the accompanying drawings in which:

5 Fig. 1 is a block diagram of an expandable telecommunications system which employs an inter-nodal network to transfer information among programmable switching nodes and a voice processing resource node which may be used in accordance with one embodiment of the present invention;

10 Fig. 2 is a block diagram of one type of programmable switching node which may be used in the system of Fig. 1;

Fig. 3 is a block diagram of the software tasks and architecture employed in the method of building a software upgrade in accordance with the present invention;

15 Fig. 4A is a schematic illustration of the matrix board in a programmable switching node including the software tasks and memory storage devices employed in the method and apparatus of downloading upgrades to the system of Fig. 1 in accordance with the present invention; and

Fig. 4B is a schematic illustration of a line card in a programmable switching node including the software tasks and memory storage devices of the line card which receive downloads of upgrades from the matrix card of Fig. 1.

### 20 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Fig. 1 shows a large capacity, expandable, fully programmable telecommunications switching system 2. Details of the construction and operation of the system 2 may be found in commonly-owned United States Patent No. 5,544,163 (Madonna), which is presently incorporated herein in its entirety by this reference.

25 To aid in the understanding of the preferred embodiment of the present invention, certain aspects of the system 2 are discussed here. The telecommunication system 2 is comprised of a host computer 4 and a plurality of inter-connected switching

- 7 -

and non-switching nodes. The embodiment illustrated in Fig. 1 includes multiple switching nodes 6a-6c. The system 2 also includes voice processing resource node 7. Any number of such non-switching nodes which perform voice processing or communication services may be used in essentially any combination with any number of switching nodes in the system 2. The nodes 6a-6c and voice processing resource node 7 are connected by an inter-nodal network 12. The inter-nodal network 12 may be a fiber optic ring. The nodes 6a-6c, as well as voice processing resource node 7, include a host interface which is connected in communicating relationship with the host 4 by a local area network (LAN) such as Ethernet or some other communication link 8. Many of the switching nodes, such as the node 6a, also have a connection to a wireline telephone service as illustrated by PSTN 14. Calls may also be connected with wireless telecommunications systems or other private networks 15.

As noted, the inter-nodal network 12 is preferably implemented using one or more fiber optic rings, 12a and 12b. However, the inter-nodal network 12 may also be implemented with any physical medium suitable for communication networks, such as, for example, wide area networks, wireless communications networks, the PSTN, (ATM/SONET) and the Internet. Using some of these other communication media to implement the inter-nodal network 12 permits the nodes to be geographically distributed over a larger area.

The overall, essentially continuous, operation of the system 2 is controlled by the host 4 as well as the software running on individual cards within the nodes 6a-6c and voice processing node 7. The host 4 is commonly implemented with a personal computer (PC) work station, fault tolerant or other computer on which the user's application runs. The host 4 communicates with the nodes 6a-6c and node 7 by exchanging messages over the LAN/link 8.

Fig. 2 illustrates the major functional components of the preferred embodiment of one type of programmable switching node 6a which may be used in the system of



- 8 -

Fig. 1. Digital or analog network/line interfaces are terminated on a series of line card input/output (I/O) cards 20. In a preferred embodiment, network/line interfaces representing a total of 2,048 ports may be terminated by line card I/O cards 20. Redundant line card I/O card 22 may also be provided, if desired.

5       A series of digital network T1, E1, J1 or analog trunk/line cards 26 communicate with line card I/O cards 20 over LC I/O lines 28. Line cards 26 are also interfaced with redundant switching buses 30a and 30b. An optional redundant line card 32 is also provided, if desired, which communicates with redundant line card I/O card 22 over redundant LC I/O bus 34. Other types of network/line interfaces such as DS3 and  
10       SONET or others may also be provided.

      A variety of communications services such as tone detection and generation, conferencing, voice-recorded announcements, call progress analysis, speech recognition, ADPCM compression and many others are provided by one or more multifunction digital signal processing (MFDSP) cards 36. ISDN Primary Rate service and other  
15       packet communications services are provided by one or more PRI-32 cards 38. Redundant MFDSP cards 36 and redundant PRI-32 cards 38 may be optionally included.

      The CPUs in the various cards exchange messages with one another over the redundant switching buses 30a and 30b. The switching buses 30a and 30b each include a high level data link control (HDLC) bus (not shown separately.)

20       A network I/O card 40a serves as an interface between inter-nodal network 12a and a CPU matrix card/nodal switch 44a. A second network I/O card 40b serves as an interface between an optional, second inter-nodal network 12b and an optional, second CPU matrix card/nodal switch 44b, which is preferably of the same construction as nodal switch 44a. The CPU/matrix card/nodal switches 44a and 44b are interfaced with  
25       switching buses 30a and 30b and with host interfaces 42a and 42b. The CPU matrix cards/nodal switches communicate with one another via midplane HDLC bus 46. If

- 9 -

desired, a redundant CPU/matrix card and a host interface (not shown) may be included.

The node 6a interfaces with the inter-nodal network 12 via network I/O card 40a having ports A and B. The network I/O card 40b provides redundancy. Interface with  
5 the inter-nodal network 12 allows the node to communicate with other switching nodes and other voice processing resource nodes in the system 2 which facilitates real time call processing functions and other services such as conferencing. The interface to the inter-nodal network 12 also provides a communication link for nodes which may not be connected directly to the host 4. The host 4 typically sends messages to the matrix card  
10 about a variety of matters such as, for example, initial configuration information for that node, system level configuration information including whether there is a change in the system due to the addition of a new node or a previously active node going out of service, or a fault in the system in which case corresponding fault isolation information is also provided.

15 The in-service upgrade method and apparatus of the present invention will now be described with reference to Figs. 3, 4A and 4B. A make/build software utility 300 and a load generator task 302 are used in an offline development system appropriate for the operation. The products of the software so developed are formed into downloadable image files which are transferred to a service provider's host computer 4 or a network  
20 file server (not shown). As will be understood by those skilled in the art, there are several ways in which the software for a base build and an upgrade may be formulated. In accordance with one preferred embodiment of the invention, the base build and the upgrade are constructed of "C" functions, i.e., C programming language source level functions. The base set of software for the system is represented in Fig. 3 by the base  
25 load files 304. The upgrade source files 306 may contain several different upgrades, 1 through n, each of which might represent an upgrade to a particular type of targeted card, such as the matrix cards.

- 10 -

The make/build utility 300 constructs a software "build" from the base load files 304 and the upgrade source files 306 in accordance with build instructions 308. More specifically, an upgrade or base build request message instruction 310 instructs the make/build utility 300 to prepare a build. The make build utility 300 employs a compiler and linker to organize and locate all the files in certain, predetermined memory portions in accordance with build instructions 308. In the first instance, the base set of files must be produced. Initially, when building the base build for the first time, there is no upgrade per se, but the upgrade source files 306 do contain information that causes the reserved memory area to be set aside. The make/build utility 300 thus creates the base build 320 with the reserved memory area 324 allotted for future upgrades. Subsequently, the base build is downloaded. The base build will be actively running on the system when a substantive upgrade (which changes code or adds additional code) is to be downloaded to the system while in service.

When preparing such an upgrade, the build instructions 308 inform the make/build utility 300 to re-build the base build even though the base build is at that time actively running on the system. Upgrade source files are edited to replace the instructions which created the reserve memory segment with the source of an actual upgrade. The make/build utility 300 is instructed via the build instructions 308 to rebuild the base build with the edited upgrade source files. The base build is rebuilt in order to assure that the upgrade is located in the reserved memory area of the architecture. Rebuilding the base set of software also establishes proper address resolution so that software routines called by the upgrade which are resident in the base build will be identified with the correct address.

The make/build utility 300 produces the upgrade build utilizing the upgrade source files 306 and the instructions 308 which cause the make/build utility 300 to place the upgrade in the reserved area 324'. The mechanism for ensuring that the upgrade is located in the reserved memory area is compiler and linker dependent. Typically, the upgrade source files contain a statement, such as a pragma statement, which

- 11 -

indicates that each of the products that comprise the upgrade are to be associated with a name. Specifically, certain products are assigned a predetermined name in accordance with a suitable naming convention. At compile time, the products, which may include code, initialized data, uninitialized data, constants and strings, each have a unique  
5 name. During link time, when all addresses get resolved and items get located in memory, the linker is given instructions by the make/build utility 300 that state that the products of the upgrade having a particular name are to be loaded into the reserved memory area 324' beginning at one particular address. Those products of the upgrade identified by a different name are to be located in the reserved memory area beginning  
10 at another address.

A map/link file is generated as a result of the build process. The map/link file 328 includes the names of the products of the upgrade as well as the addresses in the reserved memory area at which they are located. Using this information in the map/link file 328, the load generator 302 locates the beginning and end addresses of each up-  
15 grade, i.e., Upg<sub>1</sub> through Upg<sub>n</sub> in the memory image of the build. In response to a message such as "generate Upg<sub>1</sub>," the load generator 302 grabs the upgrade information after reading the addresses from the map/link file 328. After obtaining this information, the load generator 302 then extracts the upgrade portion of the build image and adds a header 332 to it so that the upgrade is ready to be downloaded onto the active system.  
20 It may perform several of the following functions depending upon the compiler and linker which are utilized in the build process: the load generator may create data records of equal length, compress the data and convert S records into binary records so that the upgrade is ready for downloading.

The header 332 indicates what the loadable image is and where it is to be loaded  
25 onto the system. More specifically, the header contains information stating which upgrade it is, 1 through n. The header 332 also includes the information about which card in the system the upgrade is to be loaded onto and where in the memory storage device on the card it is to be loaded. The header 332 further preferably includes the revision

- 12 -

identification (ID) of the base build with which the upgrade is built. This ID relates to system integrity in that it allows the card resident loader to determine whether or not the upgrade it is receiving is compatible with the base load it is running.

It may also be desirable from time to time to create a new loadable base image.  
5 A loadable base image is generated by the load generator in a similar manner as is followed for an upgrade, and is shown in Fig. 3 as base build 334 having header 336.

Now, the loadable upgrade image is ready to be downloaded onto the system. The download of upgrades is illustrated in greater detail in Figures 4A and 4B. Figure 4A illustrates the software architecture components and memory storage devices which  
10 reside on an active matrix card 400. The active matrix card 400 provides the central control for upgrade downloads. It also provides the central battery-backed repository for downloaded upgrade images. The active matrix card 400 also participates in the distribution of those upgrade images to the affected line cards and function cards.

It is noted that it may be preferable in certain circumstances to initially download  
15 a full, new base load, which replaces much of the system software, to a standby matrix card. In this way, it can be determined whether the download of a full, new base load has been successful. If the download is found to be successful, the standby matrix card can then be instructed to begin operating as the active matrix card (with the new base load running on it). The former active matrix card then is instructed to operate as  
20 the standby card, after the new software has similarly been downloaded to it.

The telecommunication system of the preferred embodiment includes a host 4 (Fig. 4A) which communicates with the inter-connected switching and non-switching nodes. A switching node includes a CPU/matrix card 400. The CPU matrix card 400 is also sometimes referred to herein as "the matrix board 400". A system loader task 402  
25 runs on the matrix card 400 and the system loader task 402 includes an upgrade manager 404. The upgrade manager 404 responds to instructions relating to upgrades, such



- 13 -

as instructions to obtain stored upgrades. The base loads 406 and the upgrades 408 can be stored on the matrix board itself, or may be stored off board on a disk or on a network file server, as desired in a particular application. Whichever way in which such files are stored, the upgrade manager 404 has access to the stored loads and can obtain  
5 the files when instructed to do so.

The matrix card 400 contains several memory storage devices associated with its CPU. The primary random access memory (RAM) 426 contains, in uncompressed form, much of the software code utilized to operate the matrix card 400. Some of the code is also stored in a read only memory (ROM) (not shown) but this code is typically  
10 software which remains unchanged after the initial configuration of the system such as its boot routine. In addition to those storage devices, a backup RAM 429 is provided which is powered by a battery for emergency use in the case of a power outage.

In order for a download to begin, the host 4 issues an upgrade download message via its Ethernet connection 8 with the matrix board 400. Preferably, the host 4  
15 downloads all upgrade information which is destined for a particular node to the active matrix card 400 of that node. In some instances, it may be desirable to download an upgrade directly to a targeted card, such as a line card. However, it is preferred to download upgrades initially to the matrix card 400, so that there is a single unit in the node that possesses a definitive list of the base loads and upgrades which reside on each  
20 card in that node. This is important for several reasons including for purposes of maintaining complete redundancy.

In response to the host-issued download message, the upgrade manager 404 running on the matrix card 400 of the node obtains the relevant upgrade 408 which may be any one of Upgrades 1 through n, from the stored load area 409. It then downloads  
25 the upgrade from the host through it (the upgrade manager 404) to the stored load area on the matrix board 400. Thereafter, as described in further detail herein, certain up-

- 14 -

grades may be distributed to another card in accordance with the instructions received by the upgrade manager 404 from the host 4.

The upgrade is loaded onto the matrix card itself if the upgrade pertains to operations of the CPU on the matrix card 400. More specifically, the upgrade information is stored in executable form in the primary RAM 426 of the matrix card 400. The upgrade manager 404 distributes or loads an upgrade 408 as shown by arrow 422 to the reserved upgrade memory segment 427 of the primary RAM 426 of the matrix board 400. As illustrated in Fig. 4A, the various upgrades, i.e.,  $Upg_1 - Upg_n$  are thus installed in the reserved upgrade memory segment 427 of the primary RAM 426. Compression of upgrades may be an option for back-up storage of the upgrade if the upgrade image is of a sufficiently large size on which to perform a compression algorithm. The identical upgrade information is downloaded to a standby matrix card (not shown) via a local matrix-to-matrix HDLC bus, in both uncompressed executable format and compressed format in a manner similar to that described with reference to the active matrix card 400.

Upgrades targeted for line or function cards are distributed by the active matrix card 400. For purposes of illustration, we have shown in Figure 4B a line card 420. It should be understood that the downloading of upgrades destined for other card types such as packet engine cards and other function cards is handled in a similar manner. Assuming that the upgrade 408 is targeted for the line cards, the appropriate sets of upgrades are distributed to the line cards 420, by the upgrade manager 404 as depicted by arrow 428 across a midplane matrix-to-line card HDLC bus 430. The upgrade is distributed across the HDLC bus 430 in response to a card upgrade download message to the upgrade manager 431 running on the line card 420.

In a manner similar to that described with reference to the matrix card 400 the upgrade manager 431 of the line card 420 then loads the upgrade as depicted by arrow

- 15 -

432 into the reserve upgrade memory segment 434 of the primary RAM 435 of the line card 420.

The line card 420 has a card monitor task 436 running on its CPU. The card monitor task 436 contains information about the tasks that are running on the card, the status of the card and whether the card is operating properly. The card monitor task 436 receives through an upgrade manager task 431 information about upgrades which have been loaded and installed to the reserve upgrade memory segment 434. Upgrades are listed and identified in the card upgrade status area 440 and the base load is identified in the card base load status area illustrated by block 444. After a card upgrade download, a list of all upgrades contained on that card is transmitted back via the HDLC bus, as depicted by arrow 448, back to the system monitor task 450 running on the matrix card 400.

When an upgrade is downloaded to the matrix card 400 or is distributed by the upgrade manager 404 of the matrix card 400, the upgrade manager 404 has certain responsibilities with respect to the upgrade. First, the upgrade manager 404 loads and stores a list of the upgrades and the base loads as shown in the card upgrade status area 454 (Fig. 4A) and the card base load status area 452 so that the system monitor task 450 has this information. The upgrade manager 404 also verifies the integrity of the upgrade. It does so by examining the header 332 (Fig. 3). The header 332 indicates whether the information contained in that image is a base load or is an upgrade. If it is an upgrade, the header identifies which upgrade it is by its revision identification number. Preferably, the revision identification number of an upgrade matches the unique revision identification number of the base build with which it is intended to operate. The upgrade manager 404 thus checks the revision number to make sure it is the correct upgrade for the base build then running on the system. It also validates the integrity of the upgrade by checking the size and the card type for which the upgrade is targeted. The upgrade manager 404 can also perform other validity checks as desired in the particular application.

- 16 -

After the matrix card 400 has received, via its system monitor 450, messages from all line cards, function cards, and other types of cards which have received downloads, it will acknowledge this back to the host 4 by sending the host 4 a per card base load status list 452 and a per card upgrade status list 454. The host 4 has the information as to which upgrades each card should contain and it compares this information with the list of upgrades received from the matrix board system monitor task 450. Any discrepancy is then resolved or an error is noted.

Once the upgrades are downloaded into the matrix cards and the appropriate upgrades are distributed to the other line and function cards in the system, and the upgrades have been verified and acknowledged to the host, the upgrades can be activated. The host 4 issues activate/deactivate messages which are passed onto an upgrade start task 448 running on the matrix card 400. In accordance with the preferred embodiment of the invention, all upgrades do not have to be activated at once. For example, all line cards of a certain type may contain a particular upgrade, however, the line cards in certain nodes, or for certain customers, may not require that upgrade at that particular time. The user may choose not to active the upgrade in every line card, but only in line cards in certain nodes.

In accordance with the present invention, activation is performed by overlaying an instruction at the beginning of an old function which transfers control to a different address that contains the new (upgrade) function. As noted, the method of activating an upgrade depends upon the processor being used and upon how the RAM is calling the routine. In one instance, a RAM may call a procedure directly by name. Alternatively, a routine may be called indirectly by reference to an address table or a jump table. In that instance, routines are called by scrolling through the table and executing the routines. In order to dynamically change the order in which the software is called, the entries in the jump table can be changed.

- 17 -

If the routine is being called by the direct approach, activation is performed in accordance with the present invention by the upgrade start task 448 which overlays a load-address-and-jump sequence 460 at the beginning of the old routine depicted by reference 462 in Fig. 4A. The number of instructions in this sequence varies depending upon the type of processor being used as will be understood by one skilled in the art. Briefly, a new address is inserted which directs the CPU to the reserve memory area 427 and thus to the upgrade Upg<sub>n</sub>. The CPU jumps to that address and executes the new (upgrade) routine. After executing the new routine, a return-to-subroutine (RTS) message 464 in the sequence returns to the same place in the RAM to which it would have returned after the old function had executed.

If the routine is being called indirectly via a jump table, the method of the present invention includes overlaying the desired entry in the jump table with the address of the new routine. The new code is executed and then the CPU is returned back to the jump table to the point where it would have been after executing the old routine.

Prior to overwriting any code, however, it is preferred to save away the portion of the old function, which is to be overlaid by activation so that the upgrade process is reversible. More specifically, the old function is saved and the address of the location of the old function is stored so that if an upgrade is to be deactivated, the jump instruction is simply removed and the old function is reinstated where it had been in the original base build.

In order to maintain the integrity of the activation operation, it is desirable that the code fragment which is being overlaid with the load-address-and jump instruction is not actually executing at the moment the activation occurs. There are several ways in which to accomplish this. One embodiment of the invention includes building the activation routine as the lowest priority software task in the computer. This means that no other task is then executing, so that code, which is then executing, will not be inadvertently overwritten. Specifically, in the illustrative embodiment, the upgrade start task



- 18 -

448 is identified as the lowest priority task on the board. The upgrade start task 448 runs when the other tasks have already run. Then, the upgrade start task 448 is given control of the processor; it locks out the other software for the short interval it takes to do the instruction overlay. Thus, the operation of overlaying an activation instruction  
5 sequence does not interfere with, or cannot be interfered with by, other software running on the computer. An alternative includes identifying which tasks are actually running at a particular point and performing the overlay if the piece of code being patched is not executing at that moment.

When the upgrade start task 448 has control of the processor, the CPU saves  
10 away the old address, overlays the new address and jump instruction so that the upgrade can be accessed. Then, interrupts are re-enabled and control of the CPU is given up. At this point, the upgrade has been validated, downloaded and activated while the overall system continues operation without interruption.

A query command can be generated by the host in which the host 4 queries the  
15 matrix card 400 for a list of upgrades that have been downloaded and whether those upgrades have been activated for each card in the node. The system monitor 450 can access the per card base load status 452 and the card upgrade status 454 and report a list of upgrades and base loads running on each card to the host 4.

In the case of a new line card being inserted or new node being added to the  
20 system, the matrix card 400 downloads the base load and then downloads all upgrades which are associated with that particular revision of the base load. Then, upgrades will be activated as desired, in accordance with the present invention.

A list of the upgrades is maintained on the battery-backed memory so that, upon a restart after a power outage or a new node being added, then the matrix card will have  
25 the necessary information about which base load and which upgrade are to be running on the system. If any piece of the software is missing then the system loader 402 ob-

- 19 -

tains this software from the stored load area, 406, 408 and it is downloaded and activated. It is also noted that configuration information is typically also downloaded when a card resets. Other information about system operations and call processing information is also downloaded in the normal course of events upon a restart.

5 Information may also be downloaded to the connect board of a voice processing node 7 (Fig. 1). The executable form of the upgrade is downloaded into its primary RAM (not shown). A standby connect board may also be provided for redundancy. In accordance with the invention, each card, whether a matrix card, a line or function card, or a connect board contains a reserved memory segment into which the upgrade image  
10 is downloaded, stored, and activated or deactivated.

It should be understood that while a preferred embodiment of the invention as used in a telecommunication system has been described herein, the invention is equally applicable to other continuously running systems which include data processing components having associated addressable memory storage devices. Such systems include,  
15 for example, air traffic control systems, RADAR systems and industrial control systems.

The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention.  
20 Therefore, it is the object of the appended claims to cover all such variations and modifications as come within the true spirit and scope of the invention.

What is claimed is:

- 20 -

## CLAIMS

- 1           1.       In a telecommunications system, comprising a plurality of intercon-  
2       nected nodes for performing telecommunications switching and other voice and data  
3       processing services, the nodes being interconnected by an inter-nodal network, and a  
4       host being connected in communicating relationship with at least one of the nodes, each  
5       node having at least one processing card containing a microprocessor with an associ-  
6       ated memory storage device, the microprocessor in the processing card running soft-  
7       ware for controlling certain functions of the card and of the node as a whole, each node  
8       also including a plurality of line cards and function cards for performing call processing  
9       functions and call connections, a method of upgrading software in at least one of the  
10      cards in a node, while the telecommunications system remains active, the method in-  
11      cluding the steps of:
- 12           (A)     establishing a reserved memory area in the memory storage device of  
13           each card targeted to receive a software upgrade;
- 14           (B)     generating a software upgrade including sets of upgrade information  
15           destined for certain predetermined targeted cards in nodes in the system, and in-  
16           cluding as part of each said set of upgrade information, instructions indicating  
17           that said set of upgrade information is to be loaded into the reserved memory  
18           area of the card for which it is targeted;
- 19           (C)     downloading said sets of upgrade information from said host to said re-  
20           served memory segments in targeted cards;
- 21           (D)     activating desired sets of upgrade information now resident on selected  
22           ones of said processing cards, line cards and function cards by overlaying an in-  
23           struction at data locations to which the processor was originally directed, said  
24           destinations being at addresses associated with software code which is to be re-  
25           placed by said software upgrade, and the jump instruction is instead read, in re-  
26           sponse to which the processor points to a new location in said reserved memory  
27           segment in which a new, upgrade code is located and is thus executed, after  
28           which the processor is instructed to return to the point to which it would have

-21-

29           returned had the prior routine been called, whereby a software upgrade is ef-  
30           fected in the system while the system continues to operate.

1           2.       The method of claim 1 including the further step of initially download-  
2           ing all software upgrade to the processing card in each node, and then distributing pre-  
3           determined sets of upgrade information from said processing cards to reserved memory  
4           segments of selected line cards and function cards in the nodes whereby a software up-  
5           grade is effected in the system while the system continues to operate.

1           3.       The method of claim 1 wherein said memory storage device includes an  
2           address table for dynamically calling up software routines, including the step of down-  
3           loading new addresses into said address table in response to which the processor trans-  
4           fers control to a different address and calls up a set of new, upgrade information corre-  
5           spondingly downloaded into the respective reserve memory segment.

1           4.       The method of claim 1 including the further step of:  
2           providing a back up memory storage device in selected cards which contains a com-  
3           pressed copy of said upgrade information and further providing said back up memory  
4           storage device with a battery.

1           5.       In a data processing system, having a base software build operating on  
2           various components of the system, with each component including an associated ad-  
3           dressable memory storage device, a method of upgrading said base software build on  
4           predetermined targeted components of the system while the system is operating, the  
5           method including the steps of:

6           (A)     establishing in the memory storage device of the a base software build a  
7           reserved memory segment and identifying the addresses of said reserved mem-  
8           ory segment;  
9           (B)     building a software upgrade to be associated with the base software  
10          build, said software upgrade being compiled and linked with the base software

-22-

- 11 build such that the addresses of instructions and portion of the code of said  
12 software upgrade correspond with those of the base software build;
- 13 (C) extracting said software upgrade from said base build and including a  
14 header in said software upgrade which includes information about the upgrade  
15 and instructions about the targeted device to which a particular portion of the  
16 software upgrade is to be downloaded;
- 17 (D) downloading said software upgrade to the reserved memory segment of  
18 the targeted device; and
- 19 (E) activating desired sections of said software upgrade while the base build  
20 is operating on said system.






FIG. 1





FIG. 2



**FIG. 3**




FIG. 4A




FIG. 4B



# INTERNATIONAL SEARCH REPORT

International Application No  
PCT/US 00/18940

## A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04Q3/545

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	HIROSHI SUNAGA ET AL: "A RELIABLE COMMUNICATION SWITCHING PLATFORM FOR QUICK SERVICE PROVISIONING" WESCON TECHNICAL PAPERS, US, WESTERN PERIODICALS CO. NORTH HOLLYWOOD, vol. 1, 18 June 1995 (1995-06-18), pages 77-82, XP000532972	1,5
Y	the whole document	3
Y	ENGRAM R L ET AL: "TRANSPARENT SOFTWARE AND HARDWARE CHANGES IN A TELECOMMUNICATION SYSTEM" INTERNATIONAL CONFERENCE ON SOFTWARE ENGINEERING FOR TELECOMMUNICATION SWITCHING SYSTEMS, GB, LONDON, I.E.E., vol. CONF. 5, 4 July 1983 (1983-07-04), pages 157-162, XP002004485 the whole document	3
-/-		

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

### \* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

7 September 2000

Date of mailing of the international search report

18/09/2000

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

Barbelanne, A

INTERNATIONAL SEARCH REPORT

International Application No PCT/US 00/18940
---

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 682 533 A (SILJESTROEMER T) 28 October 1997 (1997-10-28) column 2, line 6-43 column 4, line 14-28 column 6, line 20-36 -----	1,5

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 00/18940

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5682533 A	28-10-1997	AU 3623995 A	19-04-1996
		CN 1158662 A	03-09-1997
		EP 0784817 A	23-07-1997
		JP 10506763 T	30-06-1998
		WO 9610319 A	04-04-1996